

Writing representations over minimal fields

S. P. Glasby and R. B. Howlett

School of Mathematics and Statistics
University of Sydney, NSW 2006, Australia

ABSTRACT. The chief aim of this paper is to describe a procedure which, given a d -dimensional absolutely irreducible matrix representation of a finite group over a finite field \mathbb{E} , produces an equivalent representation such that all matrix entries lie in a subfield \mathbb{F} of \mathbb{E} which is as small as possible. The algorithm relies on a matrix version of Hilbert's Theorem 90, and is probabilistic with expected running time $O(|\mathbb{E} : \mathbb{F}|d^3)$ when $|\mathbb{F}|$ is bounded. Using similar methods we then describe an algorithm which takes as input a prime number and a power-conjugate presentation for a finite soluble group, and as output produces a full set of absolutely irreducible representations of the group over fields whose characteristic is the specified prime, each representation being written over its minimal field.

1. The main algorithm

Let $\rho: G \rightarrow \mathrm{GL}(d, \mathbb{E})$ be an absolutely irreducible representation of the group G . It is clear that there exists a subfield \mathbb{F} of \mathbb{E} , minimal with respect to inclusion, such that there exists a representation $G \rightarrow \mathrm{GL}(d, \mathbb{F})$ equivalent to ρ . If \mathbb{E} has nonzero characteristic, then \mathbb{F} is determined by ρ , and coincides with the subfield generated by the character values of ρ (see [2, VII Theorem 1.17]). Indeed, the arguments presented here yield a proof of this fact. If \mathbb{E} has characteristic zero, there may be more than one choice for \mathbb{F} .

Suppose that \mathbb{F} is a subfield of \mathbb{E} such that \mathbb{E} is a finite Galois extension of \mathbb{F} whose Galois group is cyclic, of order t , and generated by α . Assume further that the norm map from \mathbb{E} to \mathbb{F} (given by $\lambda \mapsto \lambda\lambda^\alpha\lambda^{\alpha^2}\cdots\lambda^{\alpha^{t-1}}$) is surjective. This hypothesis certainly holds if $|\mathbb{E}|$ is finite, and this is the case of principal interest to us. Our first objective is to describe a procedure which determines whether an absolutely irreducible representation $\rho: G \rightarrow \mathrm{GL}(d, \mathbb{E})$ of a finitely generated group G is equivalent to a representation $G \rightarrow \mathrm{GL}(d, \mathbb{F})$, and if so, finds an $A \in \mathrm{GL}(d, \mathbb{E})$ such that $A^{-1}\rho(g)A \in \mathrm{GL}(d, \mathbb{F})$ for all $g \in G$. Note that if g_1, g_2, \dots, g_n generate G , this condition is equivalent to $A^{-1}\rho(g_i)A \in \mathrm{GL}(d, \mathbb{F})$ for all $i \in \{1, 2, \dots, n\}$.

A basic step in our algorithm involves testing whether two given matrix representations of G are equivalent, and if they are, finding a nonsingular intertwining matrix. The naive approach to this problem involves solving nd^2 homogeneous linear equations in d^2 unknowns over the field \mathbb{E} . Computationally, this has cost $O(nd^6)$. Alternatively, there is a probabilistic algorithm, described by Holt and Rees in [1], which has expected running time $O(d^3)$. (This complexity result, and those throughout this section, assume that the cost of field arithmetic, including applying a field automorphism, is $O(1)$.)

With the notation as above, suppose that $A \in \mathrm{GL}(d, \mathbb{E})$ has the property that $A^{-1}\rho(g)A \in \mathrm{GL}(d, \mathbb{F})$ for all $g \in G$. The automorphism α of \mathbb{E} gives rise to an automorphism of $\mathrm{Mat}(d, \mathbb{E})$ (the algebra of $d \times d$ matrices over \mathbb{E}) which we also denote by α . Since the fixed subfield of α is \mathbb{F} , it is clear that $B \in \mathrm{Mat}(d, \mathbb{E})$ satisfies $B^\alpha = B$ if and only if $B \in \mathrm{Mat}(d, \mathbb{F})$. So $(A^{-1}\rho(g)A)^\alpha = A^{-1}\rho(g)A$ for all $g \in G$, and thus $C = A(A^\alpha)^{-1}$ satisfies

$$(1) \quad C^{-1}\rho(g)C = \rho(g)^\alpha \quad (\text{for all } g \in G).$$

Since ρ is absolutely irreducible, equation (1) determines C up to a nonzero scalar multiple. The first step in our procedure is, therefore, to use an algorithm such as in [1] to find (if possible) a $C \in \mathrm{GL}(d, \mathbb{E})$ satisfying (1). If no such C exists, then ρ cannot be written over \mathbb{F} ; so assume henceforth that such a C has been found.

PROPOSITION (1.1). *If $C \in \mathrm{GL}(d, \mathbb{E})$ satisfies (1), then $CC^\alpha C^{\alpha^2} \cdots C^{\alpha^{t-1}}$ equals μI where $\mu \in \mathbb{F}$ and I is the $d \times d$ identity matrix.*

Proof. Since $CC^\alpha C^{\alpha^2} \dots C^{\alpha^{t-1}}$ conjugates $\rho(g)$ to $\rho(g)^{\alpha^t} = \rho(g)$ for all $g \in G$, it must equal μI for some $\mu \in \mathbb{E}$, since ρ is assumed to be absolutely irreducible. However,

$$\mu^\alpha I = C(\mu I)^\alpha C^{-1} = C(C^\alpha C^{\alpha^2} C^{\alpha^3} \dots C^{\alpha^t})C^{-1} = CC^\alpha C^{\alpha^2} \dots C^{\alpha^{t-1}} = \mu I,$$

and so $\mu \in \mathbb{F}$, as desired. \square

The computation of μ can be effected by $t - 1$ vector by matrix multiplications, since if v is the first row of C then μ is the first component of the row vector $vC^\alpha C^{\alpha^2} \dots C^{\alpha^{t-1}}$. This has cost $O(td^2)$. If t is large compared with d , then μ may be computed at cost $O((\log t)d^3)$ by using the fact that $C_{2i} = C_i(C_i)^{\alpha^i}$ for each i , where $C_i = CC^\alpha \dots C^{\alpha^{i-1}}$.

Since the norm map from \mathbb{E} to \mathbb{F} is assumed to be surjective, there exists a $\nu \in \mathbb{E}$ whose norm is μ . We do not address here the practical problem of finding ν given μ . The methods used for storing field elements and performing field computations obviously affect this issue. (When $|\mathbb{F}|$ is bounded, there is an $O(1)$ probabilistic algorithm for computing ν .) Once ν has been found we may replace C by $\nu^{-1}C$, and assume thereafter that $CC^\alpha \dots C^{\alpha^{t-1}} = I$.

LEMMA (1.2). *If $C \in GL(d, \mathbb{E})$ satisfies $CC^\alpha \dots C^{\alpha^{t-1}} = I$, then there exists a nonzero column vector $v \in \mathbb{E}^d$ such that $Cv^\alpha = v$.*

Proof. Let $u_0 \in \mathbb{E}^d$ be nonzero, and for $i > 0$ define u_i recursively by $u_i = Cu_{i-1}^\alpha$. Observe that $u_t = u_0$. Now since the field automorphisms $\alpha^0, \alpha^1, \dots, \alpha^{t-1}$ are distinct they are linearly independent, and since the u_i are nonzero it follows that there exists a $\lambda \in \mathbb{E}$ such that $v = \sum_{i=0}^{t-1} \lambda^{\alpha^i} u_i \neq 0$. Moreover, $Cv^\alpha = \sum_{i=1}^t \lambda^{\alpha^i} Cu_{i-1}^\alpha = v$, as desired. \square

The following proposition may be viewed as a generalization of the multiplicative form of Hilbert's Theorem 90. The corresponding generalization of the additive form is trivially true.

PROPOSITION (1.3). *If $C \in GL(d, \mathbb{E})$ satisfies $CC^\alpha \dots C^{\alpha^{t-1}} = I$, then there exists an $A \in GL(d, \mathbb{E})$ with $C = A(A^\alpha)^{-1}$.*

Proof. The result is true when $d = 1$ by the multiplicative form of Hilbert's Theorem 90. Proceeding by induction, assume that $d > 1$. By Lemma (1.2)

there exists a nonzero vector v such that $Cv^\alpha = v$, and if B is an invertible matrix with v as its first column then

$$B^{-1}CB^\alpha = \begin{pmatrix} 1 & u \\ 0 & C_1 \end{pmatrix}$$

where $C_1 \in \text{GL}(d-1, \mathbb{E})$ satisfies $C_1 C_1^\alpha \cdots C_1^{\alpha^{t-1}} = I$. By the inductive hypothesis, there exists an $A_1 \in \text{GL}(d-1, \mathbb{E})$ such that $C_1 = A_1(A_1^\alpha)^{-1}$, and it follows that

$$\begin{pmatrix} 1 & 0 \\ 0 & A_1 \end{pmatrix}^{-1} B^{-1}CB^\alpha \begin{pmatrix} 1 & 0 \\ 0 & A_1 \end{pmatrix}^\alpha = \begin{pmatrix} 1 & u_1 \\ 0 & I \end{pmatrix}$$

where $u_1 = u(A_1^{-1})^\alpha$ satisfies $\sum_{i=0}^{t-1} u_1^{\alpha^i} = 0$. It follows from the additive form of Hilbert's Theorem 90 that there exists a row vector u_2 with $u_1 = u_2 - u_2^\alpha$, and then

$$A = B \begin{pmatrix} 1 & 0 \\ 0 & A_1 \end{pmatrix} \begin{pmatrix} 1 & u_2 \\ 0 & I \end{pmatrix}$$

has the required property $C = A(A^\alpha)^{-1}$. □

Note that if $C = A(A^\alpha)^{-1}$ then the map $\text{Mat}(d, \mathbb{E}) \rightarrow \text{Mat}(d, \mathbb{E})$ given by

$$\begin{aligned} X &\mapsto X + CX^\alpha + CC^\alpha X^{\alpha^2} + \cdots + CC^\alpha \cdots C^{\alpha^{t-2}} X^{\alpha^{t-1}} \\ &= A(A^{-1}X + (A^{-1}X)^\alpha + \cdots + (A^{-1}X)^{\alpha^{t-1}}) \end{aligned}$$

has image consisting of all matrices of the form AY with $Y \in \text{Mat}(d, \mathbb{F})$. These are exactly the matrices $A' \in \text{Mat}(d, \mathbb{E})$ such that $(A^{-1}A')^\alpha = A^{-1}A'$, or equivalently, $C(A')^\alpha = A'$. If X is chosen arbitrarily and $X \mapsto AY = A'$, then the probability that Y is invertible (so that $C = A'((A')^\alpha)^{-1}$) is $|\text{GL}(d, \mathbb{F})|/|\text{Mat}(d, \mathbb{F})|$. It follows that a reasonable procedure for finding an A satisfying the equation $C = A(A^\alpha)^{-1}$ is to choose $X \in \text{Mat}(d, \mathbb{E})$ randomly and compute $A = X + CX^\alpha + CC^\alpha X^{\alpha^2} + \cdots + CC^\alpha \cdots C^{\alpha^{t-2}} X^{\alpha^{t-1}}$, repeating if necessary until an invertible A is found. (One may show that $1 - |\mathbb{F}|^{-1} \geq |\text{GL}(d, \mathbb{F})|/|\text{Mat}(d, \mathbb{F})| > 1 - |\mathbb{F}|^{-1} - |\mathbb{F}|^{-2} \geq 1/4$.)

Observe that $C = A(A^\alpha)^{-1}$ combines with equation (1) to give

$$A^{-1}\rho(g)A = (A^{-1}\rho(g)A)^\alpha \quad (\text{for all } g \in G).$$

It follows that $A^{-1}\rho(g)A \in \text{GL}(d, \mathbb{F})$ for each g , and we have achieved our goal of constructing a representation equivalent to ρ with image contained in $\text{GL}(d, \mathbb{F})$. Note that if $A_i = X + CX^\alpha + CC^\alpha X^{\alpha^2} + \dots + CC^\alpha \dots C^{\alpha^{i-2}} X^{\alpha^{i-1}}$ then $A_{i+1} = X + CA_i^\alpha$, and it follows that A_t can be evaluated with $t - 1$ matrix multiplications and $t - 1$ matrix additions. It can be seen, therefore, that our procedure has expected running time $O(|\mathbb{E} : \mathbb{F}|d^3)$.

2. Absolutely irreducible representations of soluble groups

Suppose that we are given a consistent power-conjugate presentation for a finite group G . That is, G is generated by g_1, g_2, \dots, g_n , where n is the composition length of G , with defining relations

$$\begin{aligned} g_i^{p_i} &= v_i & (1 \leq i \leq n) \\ g_i^{-1} g_j g_i &= w_{ij} & (1 \leq i < j \leq n) \end{aligned}$$

where each p_i is a prime and each v_i is a word in the generators g_j for $i < j \leq n$, and each w_{ij} is a word in the g_k for $i < k \leq n$. It is clear that a group has such a presentation if and only if it is finite and soluble. Specifically, if G_i is the subgroup of G generated by g_i, g_{i+1}, \dots, g_n , then

$$(*) \quad G = G_1 \geq G_2 \geq \dots \geq G_n \geq G_{n+1} = \{1\}$$

is a subnormal series, and for each i the quotient G_i/G_{i+1} has order dividing p_i . Given that n is the composition length of G , it follows that $(*)$ is a composition series and the order of G_i/G_{i+1} is exactly p_i . We will show how the natural algorithm for constructing the absolutely irreducible representations of G (in a fixed nonzero characteristic), by working up the composition series $(*)$, can be readily adapted to ensure that each representation is written over its minimal field. We consider that we have constructed a representation of the group G_i once we have computed matrices representing the generators g_i, g_{i+1}, \dots, g_n .

For ease of exposition we let \mathbb{K} be a fixed algebraic closure of a field of prime order, and deal henceforth only with subfields of \mathbb{K} . Assume, inductively, that we have constructed representations $\sigma_1, \sigma_2, \dots, \sigma_s$ of the group

G_2 such that

- (i) each σ_i is absolutely irreducible and written over its (unique) minimal subfield of \mathbb{K} , and
- (ii) every absolutely irreducible representation of G_2 over \mathbb{K} is equivalent to exactly one of the σ_i .

Henceforth, to simplify the notation, we write $H = G_2$, $a = g_1$ and $p = p_1$.

The absolutely irreducible \mathbb{K} -representations of H are permuted by G via

$$\sigma^g(h) = \sigma(ghg^{-1})$$

for all $h \in H$ and $g \in G$. The first step is to find, for each i , which of the representations $\sigma_1, \sigma_2, \dots, \sigma_s$ is equivalent to the representation σ_i^a . If σ_i^a is equivalent to σ_i , then there exists a representation of G extending σ_i ; the minimal field for any such extension will be an extension of the field of σ_i . If σ_i^a is not equivalent to σ_i , then σ_i will be G -conjugate to $p = |G : H|$ of the representations σ_k . In this case the representation of G induced from σ_i is absolutely irreducible; however, its minimal field may be smaller than that of σ_i . Since G -conjugate representations of H yield equivalent induced representations of G , one representative only should be chosen from each G -conjugacy class.

CASE 1. Assume that \mathbb{E} is a finite field, and $\sigma: H \rightarrow GL(d, \mathbb{E})$ is an absolutely irreducible representation, with minimal field \mathbb{E} , such that σ^a is equivalent to σ .

Compute a matrix $A \in GL(d, \mathbb{E})$ such that $A\sigma(h)A^{-1} = \sigma(aha^{-1})$ for all $h \in H$. As σ is absolutely irreducible and $a^p \in H$, so $A^p = \mu\sigma(a^p)$ for some μ in \mathbb{E}^\times (the multiplicative group of \mathbb{E}). If the characteristic of \mathbb{E} equals p , then μ has a unique p th root $\nu \in \mathbb{E}^\times$. Indeed, ν is a power of μ since p is coprime to $|\mathbb{E}^\times|$. In this case there is a unique representation ρ of G extending σ , given by $\rho(a) = \nu^{-1}A$ and $\rho(h) = \sigma(h)$ for all $h \in H$. Suppose alternatively that the characteristic of \mathbb{E} is not p . In this case $\nu^p = \mu$ has exactly p solutions ν_1, \dots, ν_p in \mathbb{K} , and correspondingly there are p pairwise inequivalent extensions ρ_1, \dots, ρ_p of σ given by defining $\rho_i(a) = \nu_i^{-1}A$. For each i , the extension field $\mathbb{E}(\nu_i)$ is the minimal field for ρ_i . If $|\mathbb{E}^\times|$ is coprime to p , then one of the solutions of $\nu^p = \mu$ lies in the field \mathbb{E} , while the remaining $p - 1$ solutions generate the same field, which is the smallest extension of

\mathbb{E} whose order is congruent to 1 modulo p . If $|\mathbb{E}^\times|$ is a multiple of p , then all solutions of $\nu^p = \mu$ generate the same extension \mathbb{E}' of \mathbb{E} . Note that $|\mathbb{E}' : \mathbb{E}|$ is 1 or p , and \mathbb{E}' is the smallest extension of \mathbb{E} whose order is congruent to 1 modulo $p|\nu|$.

CASE 2. Assume that \mathbb{E} is a finite field, and $\sigma: H \rightarrow GL(d, \mathbb{E})$ is an absolutely irreducible representation, with minimal field \mathbb{E} , such that σ^a is not equivalent to σ .

Let k be the degree of \mathbb{E} over its prime subfield. If k is not a multiple of p , then \mathbb{E} is the minimal field for the induced representation σ^G . If k is a multiple of p , then \mathbb{E} has an automorphism α of order p whose fixed subfield, \mathbb{F} , is uniquely defined by $|\mathbb{E} : \mathbb{F}| = p$. In this case, if the representation $\sigma^\alpha: h \mapsto \sigma(h)^\alpha$ is not equivalent to one of the G -conjugates of σ , then \mathbb{E} is still the minimal field for σ^G ; however, if σ^α is equivalent to a G -conjugate of σ then one can readily show that σ^G is equivalent to $(\sigma^G)^\alpha$, and so the minimal field of σ^G is \mathbb{F} .

We present an explicit construction for an \mathbb{F} -representation equivalent to σ^G in the case that σ^α is equivalent to a G -conjugate of σ . Replacing α by a power of itself, we may assume that σ^α is equivalent to σ^a . Find an $A \in GL(d, \mathbb{E})$ such that

$$(2) \quad A\sigma(h)^\alpha A^{-1} = \sigma(aha^{-1}) \quad (\text{for all } h \in H),$$

and note that, by absolute irreducibility, $AA^\alpha \cdots A^{\alpha^{p-1}} = \mu\sigma(a^p)$ for some $\mu \in \mathbb{E}$. As in Proposition (1.1) we see that $\mu \in \mathbb{F}$, since

$$\begin{aligned} \mu^\alpha \sigma(a^p)^\alpha &= A^\alpha A^{\alpha^2} \cdots A^{\alpha^p} \\ &= A^{-1}(AA^\alpha A^{\alpha^2} \cdots A^{\alpha^{p-1}})A \\ &= \mu(A^{-1}\sigma(a^p)A) \\ &= \mu(A^{-1}\sigma(aa^p a^{-1})A) \\ &= \mu\sigma(a^p)^\alpha, \end{aligned}$$

where the last step follows from (2). Hence replacing A by $\nu^{-1}A$, where $\nu \in \mathbb{E}^\times$ satisfies $\nu\nu^\alpha \cdots \nu^{\alpha^{p-1}} = \mu$, we may assume that

$$(3) \quad AA^\alpha \cdots A^{\alpha^{p-1}} = \sigma(a^p).$$

The regular representation of \mathbb{E} considered as an \mathbb{F} -algebra yields an \mathbb{F} -algebra monomorphism $\phi: \mathbb{E} \rightarrow \text{Mat}(p, \mathbb{F})$, and since α is an \mathbb{F} -automorphism of \mathbb{E} there is an $M \in \text{GL}(p, \mathbb{F})$ satisfying $M^p = I$ and

$$M^{-1}\phi(\lambda)M = \phi(\lambda^\alpha) \quad (\text{for all } \lambda \in \mathbb{E}).$$

(We remark that computing ϕ and M is best done when the elements of \mathbb{E} are represented as polynomials over \mathbb{F} modulo an irreducible polynomial. In this case, the assumption in Section 1, that field arithmetic in \mathbb{E} can be performed in constant time, does not hold.) Let $\Phi: \text{Mat}(d, \mathbb{E}) \rightarrow \text{Mat}(pd, \mathbb{F})$ be defined by $\Phi((\lambda_{i,j})) = (\phi(\lambda_{i,j}))$, and define $S \in \text{GL}(pd, \mathbb{F})$ to be the diagonal sum of d copies of M . Then Φ is an \mathbb{F} -algebra monomorphism, and

$$(4) \quad S^{-1}\Phi(X)S = \Phi(X^\alpha) \quad (\text{for all } X \in \text{Mat}(d, \mathbb{E})).$$

It now follows that there is a representation $\rho: G \rightarrow \text{GL}(pd, \mathbb{F})$ such that $\rho(a) = \Phi(A)S^{-1}$ and $\rho(h) = \Phi(\sigma(h))$ for all $h \in H$, since

$$\begin{aligned} \rho(a)^p &= (\Phi(A)S^{-1})^p \\ &= \Phi(A)(S^{-1}\Phi(A)S) \cdots (S^{-(p-1)}\Phi(A)S^{p-1})S^{-p} \\ &= \Phi(A)\Phi(A^\alpha) \cdots \Phi(A^{\alpha^{p-1}}) && (\text{using (4) and } S^p = I) \\ &= \Phi(\sigma(a^p)) && (\text{by (3)}) \\ &= \rho(a^p) \end{aligned}$$

and

$$\begin{aligned} \rho(a)\rho(h)\rho(a)^{-1} &= \Phi(A)S^{-1}\Phi(\sigma(h))S\Phi(A)^{-1} \\ &= \Phi(A)\Phi(\sigma(h)^\alpha)\Phi(A^{-1}) && (\text{by (4)}) \\ &= \Phi(\sigma(aha^{-1})) && (\text{by (2)}) \\ &= \rho(aha^{-1}). \end{aligned}$$

It remains to check that ρ is equivalent to σ^G . It is clear that there exists a $T \in \text{GL}(p, \mathbb{E})$ such that

$$T\phi(\lambda)T^{-1} = \text{diag}(\lambda, \lambda^\alpha, \dots, \lambda^{\alpha^{p-1}})$$

for all $\lambda \in \mathbb{E}$. Furthermore, if v_i denotes the $(i+1)$ th row of T and V_i denotes the subspace of \mathbb{E}^{pd} comprising the elements of the form $(\lambda_1 v_i, \lambda_2 v_i, \dots, \lambda_d v_i)$ where $\lambda_1, \lambda_2, \dots, \lambda_d \in \mathbb{E}$, then

- (i) $\mathbb{E}^{pd} = V_0 \oplus V_1 \oplus \dots \oplus V_{p-1}$,
- (ii) each V_i is $\rho(H)$ -invariant, inducing an action equivalent to σ^{a^i} , and
- (iii) $V_i \rho(a) = V_{i+1}$, where the subscripts are read modulo p .

Note that (ii) follows from $v_i \phi(\lambda) = \lambda^{a^i} v_i$, and (iii) follows from the equation $\rho(a)\rho(h)\rho(a)^{-1} = \rho(aha^{-1})$. These conditions guarantee that ρ is equivalent to σ^G , as required. We have thus achieved our goal of constructing the absolutely irreducible representations of G over their minimal fields.

References

- [1] Derek F. Holt and Sarah Rees, *Testing modules for irreducibility*, J. Aust. Math. Soc. (A) **57** (1994), 1–16.
- [2] B. Huppert and N. Blackburn, *Finite Groups II*, Springer-Verlag, Berlin, 1982.